




# An array of opportunity



**D**espite its growing popularity, the programmable logic and fpga market is a tough place to do business. So it's a brave company that throws its hat into that ring. But that's what MathStar has done. Not only to get a toe hold in the market, but to do so on the back of a completely different approach.

An indication of the scale of the problem comes from a quick look at the company's history; it was founded in 1997. Since then, it has been developing what it calls field programmable object arrays, or fpoas. These devices are said to be capable of running at clock rates of up to 1GHz – at least twice the speed of current leading fpga architectures. Company president and ceo Doug Pihl claimed the approach represented the next generation of programmable logic. “But it offers a higher level of performance,” he continued.

So what is an fpoa? Essentially, it's a collection of silicon objects in the form of arithmetic logic units (ALU), register files (RF) and multiply accumulators (MAC) that are linked using an interconnect fabric running at 1GHz (see figure 1).

Will the ‘David’ fpoa fell the ‘Goliath’ that is the fpga?

By **Graham Pitcher**.

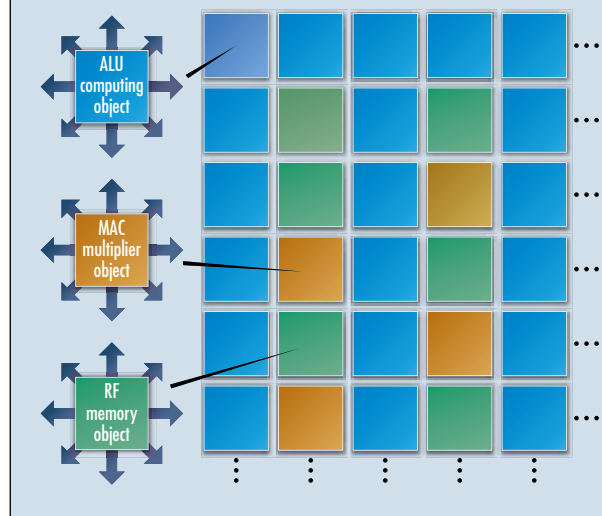
Each silicon object is a 16bit processing element which can be interconnected to provide massively parallel computing devices. Silicon objects are programmable individually and can act autonomously.

Already, MathStar is into its second generation of fpoas – currently being marketed under the Arrix brand. “The first generation worked well enough for us to demonstrate performance,” Pihl continued. “Since then, we’ve been working to productise the device.”

MathStar's approach to the market is based, predominantly, on one premise – that its technology allows the user to access levels of performance which haven't previously been achievable using programmable logic. “Our potential customers,” Pihl contended, “are struggling to get the performance levels they want from fpgas.”

He described the early days of MathStar. “When we began development, we wanted to make programmable devices, but incremental change wasn't good enough; we wanted to create a new architecture.” What it came up with was the silicon objects concept. “They are com-

**Figure 1: The FPOA architecture**





“Even if you use all the resources available on the device, it runs at 1GHz.”

Doug Pihl, **MathStar**

puting elements,” Pihl explained, “which can be used to create arrays. The array can then be programmed at a functional level, rather than at the gate level, as is the case with fpgas.”

The array is created using an interconnect ring. Each silicon object is surrounded by eight other objects – which can be ALUs, RFs or MACs. Pihl said: “This gives the designer the ability to create arrays of differ-

ent sizes and mixes to tailor to different market segments.” One of the first Arrix chips features a 20 x 20 array of silicon objects.

According to Pihl, the interconnect ring brings a number of advantages, including high performance bandwidth between objects. “It also brings flexibility,” he continued. “Programming at the gate level is too limiting and is increasingly bringing significant timing problems.”

At first sight, a chip running at 1GHz is likely to make the timing issue even trickier, but Pihl disagrees. “We want to get rid of timing problems. Whilst the chip operates at 1GHz, designers don’t have to worry about timing closure; they don’t have the gate level timing closure issues of the programmable logic world.”

The interconnect architecture is a patented part of MathStar’s offering. Each silicon object can make 18 connections to other objects in the array (see figure 2). These links are 21bit wide and run at 1GHz. Interconnect is on two levels. There are eight full duplex connections to surrounding objects through ‘Nearest Neighbour’ links. These are rectilinear and diagonal, with a limit of one object in one clock cycle.

Then 10 further connections allow communication with more distant objects. This approach – Party Line Interconnect – allows links with objects up to four steps away in the array in one clock cycle.

In Pihl’s opinion, using fpgas incurs a time to market penalty – an interesting concept, bearing in mind that programmable logic companies make a similar claim when comparing their technology to asics. He claims that a high performance fpoa based design can be completed in little more time than that taken for a low performance fpga design.

The reason, Pihl explained, was timing. “If you try to move your clock speed up in an fpga, you’ll have problems with timing closure. Our chip runs at 1GHz and always runs at 1GHz. Even if you use all the resources available on the device, it runs at 1GHz.

“Once you’ve designed the system and distributed the code amongst the objects, it still runs at 1GHz. The design tools for Arrix are cycle accurate, so you know that, once you go through the design process, what the timing will be.”

Pihl outlined the differences between fpoas, fpgas and dsp. “FPGAs have many parallel resources, slow clock speeds and require timing closure. DSPs have few parallel resources, fast clock speeds, but you have little visibility into timing. Our fpoa, by contrast, offers 400 parallel resource, a 1GHz clock speed and no timing closure.” And it is this contrast which, he hopes, will attract current fpga users. “They represent 80% of our target customers,” he claimed.

The design process starts in Mentor Graphics’ Visual Elite software, which allows design creation and editing, test bench creation and SystemC simulation. Once this has been done, the design moves on to what MathStar calls ‘coasting’. COAST – the connection and assignment tool – boasts what Pihl calls ‘rudimentary automation’. “It accelerates floorplanning,” he noted, “and once you’ve finished coasting, an object compiler generates the image.”

He doesn’t see a problem with the lack of design automation. “Our configuration sizes are probably only 10% of an fpga size and this gives the opportunity to reconfigure dynamically, which would take far too long in an fpga.”

The first Arrix family member is the MOA2400D, with a 20 x 20 array. This breaks down into 256 ALUs, 80 RFs and 64 MACs. A bandwidth of 64Gbit/s is claimed and the device has two high speed lvds ports. There is 228kbit of on chip sram and the part can access up to 288Mbyte of off chip rldram. Currently produced by TSMC on a 130nm process, the chip comes in a 31 x 31mm bga and a development board is available.

Initial target markets are machine vision and professional video. “But we’re working on a 90nm version,” Pihl observed, “which will allow us to extend these markets to include test and measurement and medical. FPOAs will be suited to any application where high performance and programmability are the goals.”

MathStar hopes this 90nm part will sample this year, followed by an upgraded device offering a 20 x 30 object array with more internal memory and ‘architectural tweaks’. The first 90nm part is likely to feature more MACs than the current device, but the later offering is likely to be ‘ALU heavy’. ☺

**Figure 2: FPOA connectivity**

